

TIM, a ray-tracing program for forbidden ray optics

Dean Lambert¹, Alasdair C. Hamilton, George Constable, Harsh Snehanhu², Sharvil Talati², Johannes Courtial*

SUPA, School of Physics & Astronomy, University of Glasgow, Glasgow G12 8QQ, United Kingdom

Abstract

TIM (The Interactive METATOY) is a ray-tracing program specifically tailored towards our research in METATOYs, which are optical components that appear to be able to create wave-optically forbidden light-ray fields. For this reason, TIM possesses features not found in other ray-tracing programs. TIM can either be used interactively or by modifying the openly available source code; in both cases, it can easily be run as an applet embedded in a web page. Here we describe the basic structure of TIM's source code and how to extend it, and we give examples of how we have used TIM in our own research.

Keywords: ray tracing; geometrical optics; METATOYs

PROGRAM SUMMARY

Manuscript Title: TIM, a ray-tracing program for forbidden ray optics

Authors: Dean Lambert, Alasdair C. Hamilton, George Constable, Harsh Snehanhu, Sharvil Talati, Johannes Courtial

Program Title: TIM

Journal Reference:

Catalogue identifier:

Licensing provisions: GNU GPL

Programming language: Java

Computer: Any computer capable of running the Java Virtual Machine (JVM) 1.6

Operating system: Any; developed under Mac OS X Version 10.6

RAM: typically 145 MB (interactive version running under Mac OS X Version 10.6)

Keywords: ray tracing, geometrical optics, METATOYs

Classification: 14 Graphics, 18 Optics

External routines/libraries: JAMA [1] (source code included)

Nature of problem:

visualisation of scenes that include scene objects that create wave-optically forbidden light-ray fields

Solution method: ray tracing

Unusual features:

specifically designed to visualise wave-optically forbidden light-ray fields; can visualise ray trajectories; can visualise geometric optic transformations; can create anaglyphs (for viewing with coloured "3D glasses") and random-dot autostereograms of the scene; integrable into web pages

Running time:

Problem-dependent; typically seconds for a simple scene

References

- [1] JAMA: A Java Matrix Package, <http://math.nist.gov/javanumerics/jama/>

1. Introduction

TIM was originally conceived as a tool to allow experimentation (although only in the computer) with novel optical components called METATOYs [1] prior to building them (TIM is an acronym for *The Interactive METATOY*). Since then, it has developed into a much more general ray-tracing program, suitable for use in optics research (including, but not limited to, METATOYs research), but also for simply playing with.

METATOYs are surfaces that appear to change the direction of light in ways that often result in wave-optically forbidden light-ray fields [2]. Of course, METATOYs cannot *actually* create wave-optically forbidden light-ray fields; what they do is create light-ray fields that are visually almost indistinguishable from the forbidden fields. Of particular interest to us are surfaces that perform a generalisation of refraction: they change the direction of transmitted light rays according to laws that can be much more general than Snell's law. In TIM, such generalised refraction can be described in terms of a complex representation, which is explained in section 2. The ability to handle very general surface properties, which enables the visualisation of scenes that include objects with METATOY surfaces (Fig. 1), is TIM's key speciality.

But TIM has other specialities, which support different aspects of our research. TIM can

- simulate photos taken with cameras that can focus on almost arbitrary surfaces (section 3);
- simulate surfaces that "teleport" light rays to corresponding positions on other surfaces, from where the light rays then continue (section 4);

*Corresponding author

E-mail address: johannes.courtial@glasgow.ac.uk

¹Now at SUPA, School of Physics & Astronomy, University of Edinburgh, Edinburgh EH9 3JZ, United Kingdom.

²HS and ST contributed while visiting from the Indian Institute of Technology Delhi, Hauz Khas, New Delhi 110 016, India.

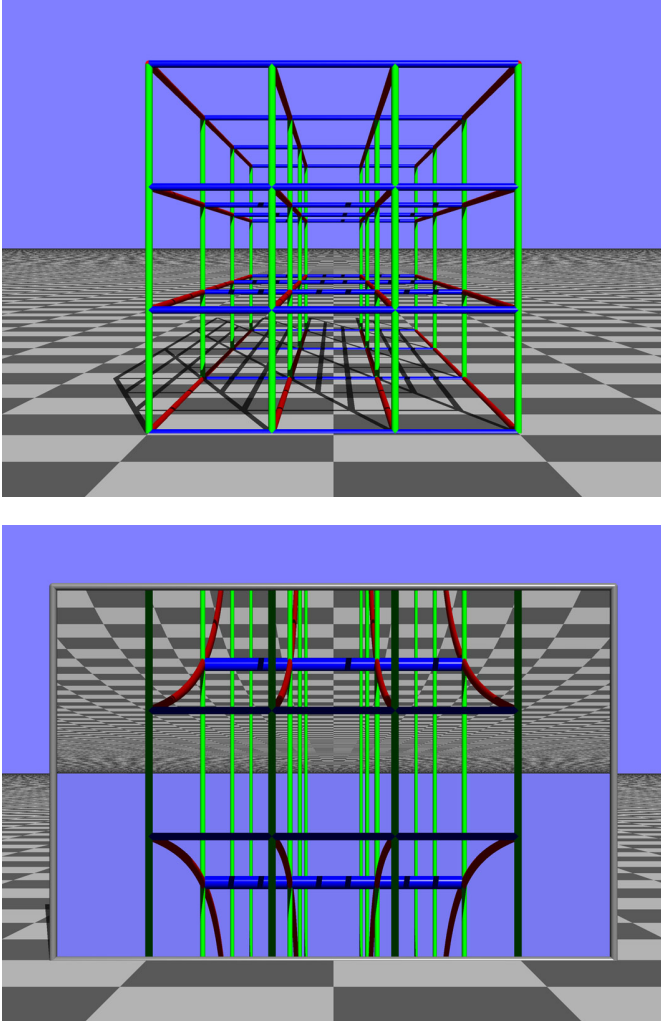


Figure 1: Simulated view of a cylinder lattice, seen on its own (top) and through a window that changes the sign of the vertical light-ray-direction components (bottom).

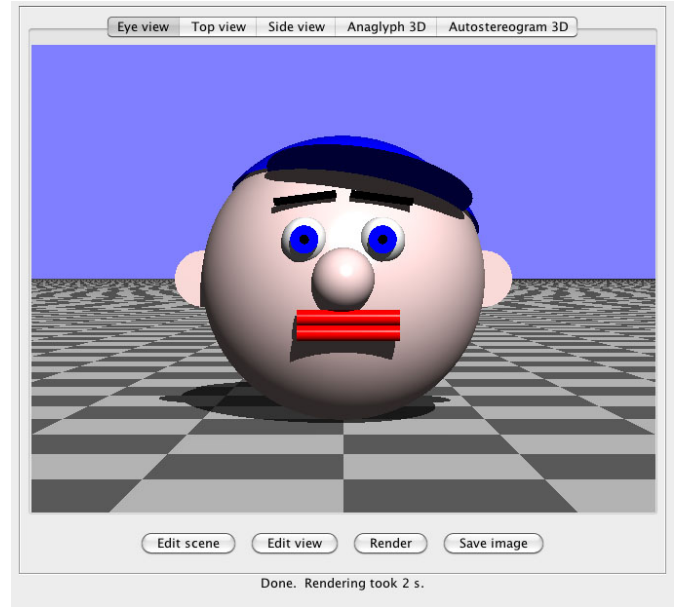


Figure 2: TIM, running as an interactive Java application on an Intel MacBook. The central image is the rendered view of the default scene. The Java applet version looks identical, apart from the “Save image” button being absent in the Java applet version because of security restrictions.

- visualise the trajectories of individual light rays (section 5);
- render scenes as anaglyphs for viewing with red/cyan anaglyph glasses (section 6);
- render scenes as random-dot autostereograms (section 7).

There is one more speciality: TIM can be run as an interactive Java applet (Fig. 2), which can easily be embedded in internet pages¹. We use this capability to disseminate our research over the internet, in a manner that invites playful exploration. The use of this interactive version of TIM is described in a user guide [3].

Sometimes our research requires capabilities which are not yet built into the interactive version of TIM. This then requires modification of the source code, and sometimes the modifications become part of the interactive version of TIM. As TIM is open-source software, in principle everybody can do this. The aim of this paper is to encourage this: to invite others to play with the interactive version of TIM, and to entice them to download and modify TIM’s source code. The paper contains several appendices aimed at facilitating the source-code modification by outlining the implementation of ray tracing, which forms the core of TIM’s source code (Appendix A); the overall structure of TIM’s source code (Appendix B); and how to perform a number of code-modification tasks, including rendering a specific scene by modifying the source code for the default non-interactive TIM Java application (Appendix C), adding a

¹An example can be found at <http://www.physics.gla.ac.uk/Optics/play/TIM/>.

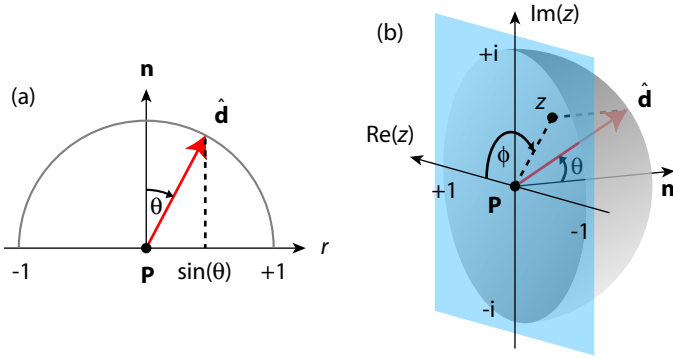


Figure 3: Representations of the normalised light-ray direction $\hat{\mathbf{d}}$ at a point \mathbf{P} on a surface. (a) Orthographic projection onto a real axis tangential to the surface at \mathbf{P} and with its origin at \mathbf{P} . The real axis lies in the plane of incidence, which is the plane through \mathbf{P} that is spanned by the local surface normal, \mathbf{n} , and the normalised light-ray direction, $\hat{\mathbf{d}}$. (b) Orthographic projection into a complex plane tangential to the surface at \mathbf{P} and with its origin at \mathbf{P} .

new class of scene object (Appendix D), and adding a new class of surface property (Appendix E).

2. Complex representation of generalised refraction

TIM exists because we wanted to see scenes that include METATOYS. Surface properties describing METATOYS therefore play a key role in TIM. TIM describes almost all METATOY surface properties in terms of a complex representation introduced in Ref. [4].

The complex representation is itself an extension of the way ray direction is represented in Snell’s law,

$$n \sin \theta = n' \sin \theta', \quad (1)$$

where θ is the angle between the incident light-ray direction and the surface normal at the point \mathbf{P} where the ray intersects the surface, θ' is the angle between the outgoing light-ray direction and the surface normal at \mathbf{P} , and n and n' are the refractive indices that meet at \mathbf{P} . Each sine of an angle with the local surface normal can be interpreted as the number at the orthographic projection of the tip of the unit vector $\hat{\mathbf{d}}$ representing the corresponding light-ray direction onto a real axis that is tangential to the surface at \mathbf{P} , has its origin located at \mathbf{P} , and is lying in the plane of incidence (the plane spanned by the incident light-ray direction and the surface normal); this is shown in Fig. 3(a). Snell’s law can then be interpreted as a simple multiplication of the projection of the light-ray direction by a factor n/n' .

We now replace the real axis with a complex plane (Argand plane), again tangential to the surface at \mathbf{P} and with its origin at \mathbf{P} . Light-ray direction can then be described in terms of the complex number z at the orthographic projection of the tip of the unit light-ray-direction vector $\hat{\mathbf{d}}$ into this complex plane (Fig. 3(b)). In this representation, Snell’s law is still described

by a simple multiplication of z by a factor n/n' . That the outgoing light ray also lies in the plane of incidence is explicitly contained in this formulation, but not in Snell’s law [4]. Rotation of the light-ray direction by an angle α around the local surface normal [5] is described by multiplication of z by a factor $\exp(i\alpha)$ [4]. Other light-ray-direction mappings can be described by other complex mappings $z \rightarrow z'(z)$; the visual effects due to such mappings are investigated in more detail elsewhere [6].

3. A camera that can focus on almost arbitrary surfaces

TIM has the ability to simulate photos taken with a camera that can focus on almost arbitrary surfaces. In this section we explain how TIM simulates such a camera.

Focussing matters only in cameras with a non-zero aperture size. (In photos taken by cameras with a point-like aperture — pinhole cameras — everything is imaged as sharply as diffractions permits.) TIM simulates a camera with a finite-size aperture by backward-tracing, starting from each pixel, a number of light rays which have passed through different points on the aperture, and averaging the resulting colours. The points on the aperture through which the backwards-traced light rays emerge are randomly chosen. Which direction these light rays have as they leave the aperture is determined by the imaging properties of the lens: the light rays originate at the position of a particular pixel, and so they have to leave the lens in the direction of the pixel’s *image* formed by the lens.

In a real lens, the images of all detector pixels lie, to a good approximation, in a plane. By allowing the images of the detector pixels to lie on a much more general surface, the *focus surface*, TIM simulates a camera which focusses on this focus surface.

The focus surface is defined as follows. The camera has associated with it a number of scene objects that define the so-called *focus scene*, a scene in addition to the scene TIM renders. The focus surface is then defined as those points in the focus scene visible to an observer placed at the centre of the aperture.

In the case of a thin lens, the image of any point, and specifically any detector pixel, lies somewhere along the continuation of the straight line from the point to the centre of the lens. We generalise this here such that, in TIM’s camera that focusses on non-planar surfaces, the image of any detector pixel lies somewhere along the continuation of the straight line from the pixel position to the centre of the aperture. Fig. 4 illustrates this.

Conveniently, the position of the image of a particular pixel at position \mathbf{P} can be found using functionality already built into TIM, namely its capability to find the closest intersection between an arbitrary light ray and a group of scene objects, one of the key capabilities for ray tracing: all that is required is finding the closest intersection between the focus scene and a ray that starts from the centre of the aperture, \mathbf{C} , with a direction given by $(\mathbf{C} - \mathbf{P})$, the direction from the detector pixel to the centre of the aperture.

Fig. 5 shows an example of a scene rendered for a non-planar focus surface. The focus scene consists of a few — but not all

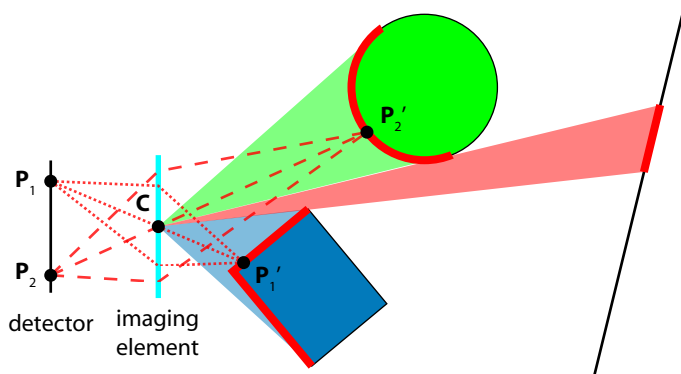


Figure 4: Construction of the focus surface and the position of the image of any detector pixel. The focus surface (thick red line) consists of those parts of the focus scene visible from the centre C of the aperture of the imaging element. The position of the image of any particular detector pixel P lies on the intersection between the focus surface and the continuation of the straight line between the pixel position and the point C . The figure shows the positions of the images of two pixels, P_1 and P_2 . In the example shown here, the focus scene consists of three objects: a circle, a rectangle, and a line.

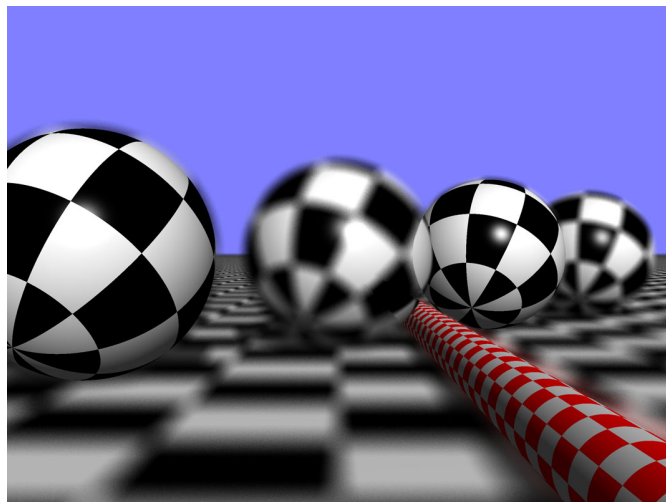


Figure 5: Example of a scene rendered with a non-planar focus surface. The focus scene consists of two of the four spheres, the chequered cylinder, and a plane in the far distance.

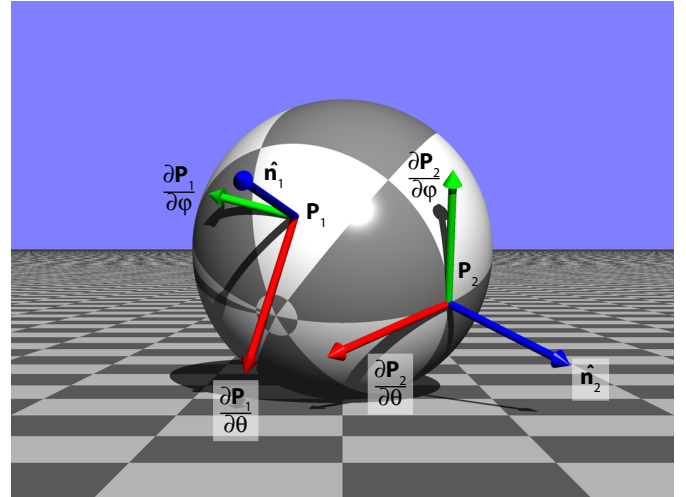


Figure 6: Parametrisation of a scene object. Each point on the surface is described by a pair of surface coordinates, in the picture the spherical coordinates θ and ϕ , defined with respect to an arbitrary zenith direction (here $(0.408, 0.408, 0.816)$) and direction of the azimuth axis (here $(0.913, -0.183, -0.365)$). This parametrisation of the surface has been indicated by covering it in a checkerboard pattern with tiles of side length 1 in both θ and ϕ . The local surface-coordinate axes, $\hat{\theta}_i = \partial \mathbf{P}_i / \partial \theta$ and $\hat{\phi}_i = \partial \mathbf{P}_i / \partial \phi$, together with the local surface normals, $\hat{\mathbf{n}}_i$, are shown for two points, \mathbf{P}_i ($i = 1, 2$). The sphere has radius 1 and is centred at $(0, 0, 10)$.

— of the scene objects in the scene, and a distant plane. The objects that are part of the focus scene can clearly be seen to be rendered sharply; those that are not are blurred.

4. Teleporting surface property

An ideal lens takes the field in one plane and creates an image of this field in another plane. The image is stretched in the transverse directions, but not otherwise distorted.

Sometimes it is desirable to create an image that is distorted according to a specific mapping between the coordinates in the object and image planes. This is called a geometrical optical transformation, and it can be approximated holographically [7]. Our own interest in geometrical optical transformations stems from the application of a polar-to-Cartesian transformation between two planes to the detection of optical angular momentum [8].

For geometrical optical transformations, coordinates are clearly important. Many of TIM's scene objects have associated with them a two-dimensional coordinate system that parametrises their surface, i.e. each point \mathbf{P} on the scene object's surface is described by a pair of associated surface coordinates, c_1 and c_2 . For example, positions on a plane are described by Cartesian surface coordinates; positions on a circular disc are described by their polar coordinates r (the distance from the centre) and ϕ (the azimuthal angle); positions on a sphere are described by their spherical polar coordinates θ (the polar angle) and ϕ (the azimuthal angle). Scene objects parametrised in this way can also calculate the local surface-coordinate axes for

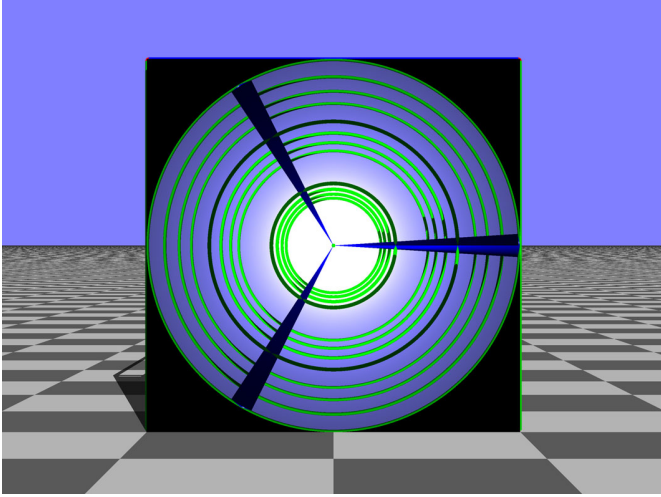


Figure 7: Simulated view through a Cartesian-to-polar converter. In the example shown here, the converter has been placed immediately in front of the cylinder lattice shown in Fig. 1. The converter distorts vertical cylinders (green) into circles centred on the origin, and horizontal cylinders (blue) into radial lines. A brightening of the central region due to area elements being transformed to a different size (see Eqn (5)) is clearly visible. The converter consists of a circular disc (the origin object), parametrised in terms of distance from the centre and azimuthal angle ϕ , with a teleporting surface whose target object is a black square immediately behind the disc, parametrised in terms of Cartesian coordinates. All coordinates are scaled to range from 0 to 1.

any point on the surface. These are the vectors

$$\hat{\mathbf{c}}_1 = \frac{\partial \mathbf{P}}{\partial c_1}, \quad \hat{\mathbf{c}}_2 = \frac{\partial \mathbf{P}}{\partial c_2}; \quad (2)$$

they respectively point in the direction in which the corresponding surface coordinate changes at the point \mathbf{P} , and their respective length indicates the distance on the surface over which the corresponding surface coordinate changes by 1. Fig. 6 shows the local surface-coordinate axes for two positions on a sphere. The surface coordinates and surface-coordinate axes, respectively, play a key role in the calculation of the starting point and direction of the continuation of the incident light ray.

TIM can simulate geometrical optical transformations using an unusual surface property. In TIM's implementation of ray tracing (see Appendix A), surface properties are responsible for returning the colour of light leaving a specific point on the surface in a specific direction. Finding this colour often requires further tracing of the incident ray, for example in the case of a specularly reflective surface, where the continuation of the ray leaves the same point on the surface with a new direction given by the law of reflection. Geometrical optical transformations can therefore be implemented in the form of a surface property that continues tracing the incident light ray, starting from a transformed position and with a suitably transformed direction.

TIM's teleporting surface property does precisely this. A teleporting surface has associated with it a destination object; note that the link is one-way. We call the object with the teleporting surface the origin object. The mapping from the position \mathbf{P} where the incident light ray intersects the origin object to the position where its continuation leaves the destina-

tion object's surface, \mathbf{P}' , is defined in terms of the surface-coordinate systems (Fig. 6) associated with the two objects' surfaces. Specifically, if the position \mathbf{P} where the incident light ray intersects the origin object's surface is described by some values of its surface coordinates, then the point \mathbf{P}' where the ray's continuation leaves the destination object's surface is the point at which the destination object's surface coordinates take those same values. For example, if a planar origin object is parametrised in terms of polar surface coordinates r and ϕ , and the destination object associated with its teleporting surface property is also planar and parametrised in terms of Cartesian surface coordinates, then this setup is a polar-to-Cartesian converter for backwards-propagating light rays, which means it is a Cartesian-to-polar converter for forwards-propagating light rays. Fig. 7 shows a three-dimensional (3D) lattice of cylinders seen through such a setup. It can clearly be seen that lines of constant z and x value (i.e. vertical lines) in the cylinder lattice become lines of constant radius (i.e. circles) when seen through the converter, and that lines of constant z and y value (horizontal lines) become lines of constant azimuthal angle (spokes).

The corresponding mapping of the light-ray direction is based on wave-optical considerations. Specifically, we assume that phase and intensity of the (scalar) optical field on the teleporting surface gets mapped onto the destination object's surface. In the ray-optics limit of wave optics [9], which is arguably appropriate here, the light-ray direction is proportional to the direction of the gradient of the phase. The wave's local phase gradient then defines the new light-ray direction. In TIM, light-ray direction is normalised and therefore naturally interpreted as normalised phase gradient.

The components of the phase gradient are, of course, the rate of change of phase in the corresponding directions. Consider the two coordinate systems that describe the surfaces of the origin object and of the destination object. Now consider the component of the incident light-ray direction in the direction of the origin object's first surface coordinate, c_1 . A value g_1 means that the phase changes locally at a rate of g_1 full 2π phase cycles over the distance on the surface in which c_1 changes by 1. On the destination object, the phase then changes locally (at the point \mathbf{P}' where the light ray continues) at a rate of g_1 full 2π phase cycles over the distance in which the destination object's first surface coordinate, c'_1 , changes by 1. The ratio of the phase gradients in the direction of the first surface coordinate, which is the ratio of the light-ray components in the direction of the first surface coordinates, is therefore given by the ratio of the distances over which the first surface coordinate changes by 1 in the origin object and in the destination object. These distances are given by length of surface-coordinate axes, $|\hat{\mathbf{c}}_1| = |\partial \mathbf{P} / \partial c_1|$ and $|\hat{\mathbf{c}}'_1| = |\partial \mathbf{P}' / \partial c'_1|$. A similar argument can be made for the second surface coordinate, c_2 . The components of the direction of the continuation of the light ray in the directions of the destination object's first and second surface coordinates, d'_1 and d'_2 , are then

$$d'_1 = \frac{|\hat{\mathbf{c}}_1|}{|\hat{\mathbf{c}}'_1|} d_1, \quad d'_2 = \frac{|\hat{\mathbf{c}}_2|}{|\hat{\mathbf{c}}'_2|} d_2, \quad (3)$$

where d_1 and d_2 are the components of the direction of the in-

cident light ray in the directions of the first and second surface coordinates of the origin object.

The component of the light-ray-direction vector in the direction of the surface normal is chosen such that the length of the light-ray-direction vector remains unchanged. This correctly represents the case of the continuation of the ray travelling in a medium with the same refractive index as the medium in which the incident ray was travelling.

One further consideration is a concomitant change in brightness. TIM assumes that no power is lost during teleportation, so the power entering an area element dA on the surface of the origin object at position P is the same as that exiting the corresponding area element dA' on the surface of the destination object at P' . The light intensities (power per area) at P and P' , I and I' , are then given by the equation $I dA = I' dA'$. The ratio of the area elements is given by

$$\frac{dA'}{dA} = \frac{|\hat{\mathbf{c}}'_1 \times \hat{\mathbf{c}}'_2|}{|\hat{\mathbf{c}}_1 \times \hat{\mathbf{c}}_2|}, \quad (4)$$

and so

$$I = I' \frac{|\hat{\mathbf{c}}'_1 \times \hat{\mathbf{c}}'_2|}{|\hat{\mathbf{c}}_1 \times \hat{\mathbf{c}}_2|}. \quad (5)$$

It is worth noting that the teleporting surface property can be used for purposes other than implementing geometrical optical transformations. For example, consider a planar origin object placed immediately in front of the camera and a planar target object placed elsewhere, both parametrised in terms of Cartesian coordinate systems. If the scale of the two coordinate systems is the same, i.e. if the surface-coordinate axes on the origin and target objects are of the same length, then the effect is a simple change of camera position and viewing direction.

5. Visualisation of light-ray trajectories

One of TIM's capabilities, namely the visualisation of light-ray trajectories, can be very helpful in understanding the effect of optical components. Fig. 8 shows a cone of light rays being traced through a window that rotates the local light-ray direction through 90° , showing, for example, that such a window would not create an image of a point light source at the apex of the cone.

TIM visualises the trajectories of specific light rays in three steps:

1. trace the light rays, keeping a list of the points where each light ray intersects a scene object;
2. for each segment of the above light-ray trajectory, i.e. between each pair of neighbouring intersection points, add a cylinder to the scene;
3. visualise the scene.

The first point uses the ray-tracing methods already built into TIM, but those methods needed to be extended slightly to keep track of ray trajectories. This requires the ability to deal with rays branching, which occurs whenever a ray encounters an object with multiple surface properties that require further ray

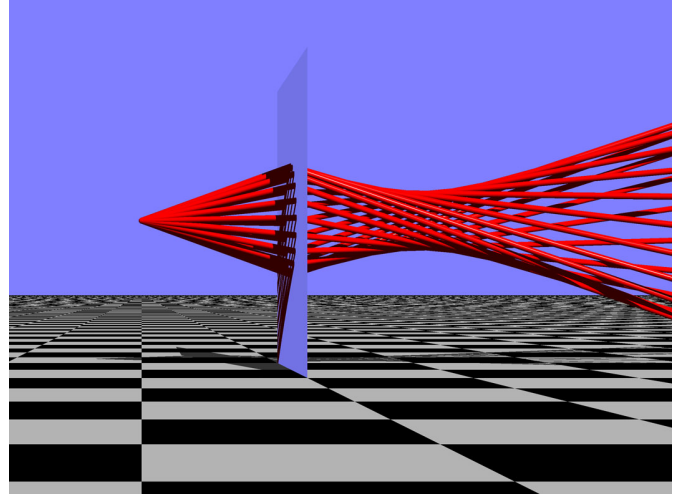


Figure 8: Visualisation of light-ray trajectories. A cone of light-ray trajectories originating from a point in front of a window that rotates the light-ray direction by 150° around the local window normal is converted into a twisted bundle of rays.

tracing, such as a partially transmissive mirror. In TIM, whenever a ray hits a surface that requires further ray tracing, a new ray is created, added to the list of the ray's branches, and traced. A light ray's full trajectory is then stored in the form of a list of positions where the main branch intersects scene objects, and the list of the branches.

6. Anaglyphs

Parallax causes a scene to look different when viewed from two different positions. This is called binocular disparity; the process of deriving depth perception from the two different views the two eyes receive is called stereopsis [10]. In anaglyph images [11], the two views are superposed, but in different colours. When viewed through suitable anaglyph glasses, i.e. a different colour filter in front of each eye which in the simplest case completely filters out the view intended for the other eye, different images can be presented to the eyes, and stereopsis can lead to depth perception.

TIM can create anaglyph images intended for viewing with red/cyan anaglyph glasses. Two images are calculated for camera positions that differ by a sideways displacement. These two images can then be turned into anaglyph images in two different ways:

1. The red component of the image is the luminosity of the left-eye image, the blue component is the luminosity of the right-eye image. The resulting anaglyph has lost colour information.
2. Following "recent simple practice" [11], the blue and green components are removed from the image that corresponds to the left eye, the red component is removed from the right-eye image, and the two images are superposed. The resulting anaglyph includes colour information, but does not work very well for objects of certain colours.

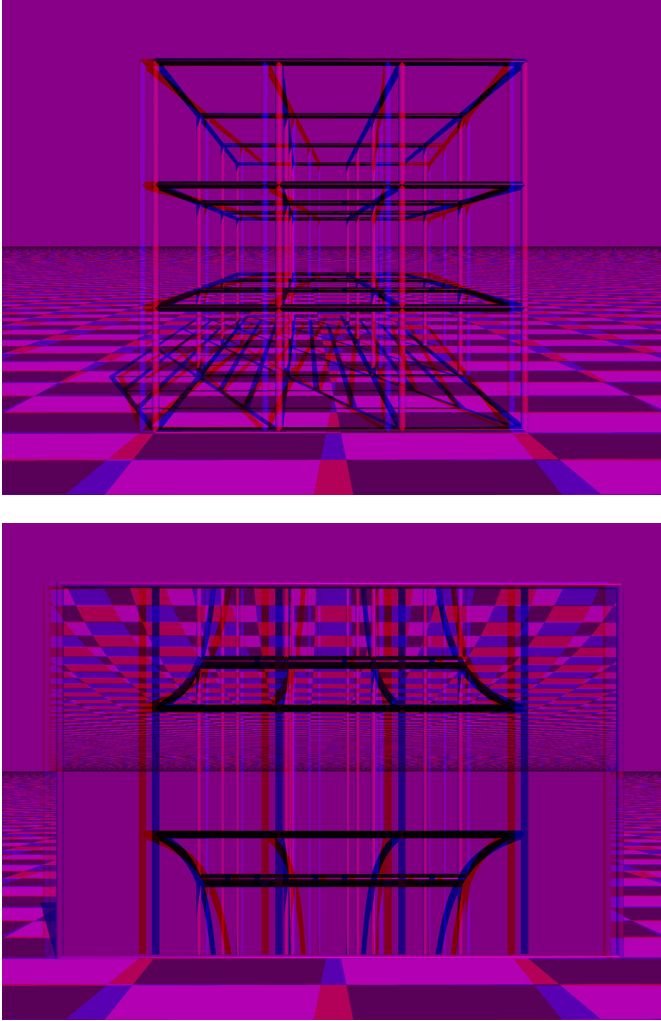


Figure 9: Anaglyph versions of the images in Fig. 1, which show a cylinder lattice *au naturel* (top) and when seen through a window that inverts the vertical ray-direction component (bottom).

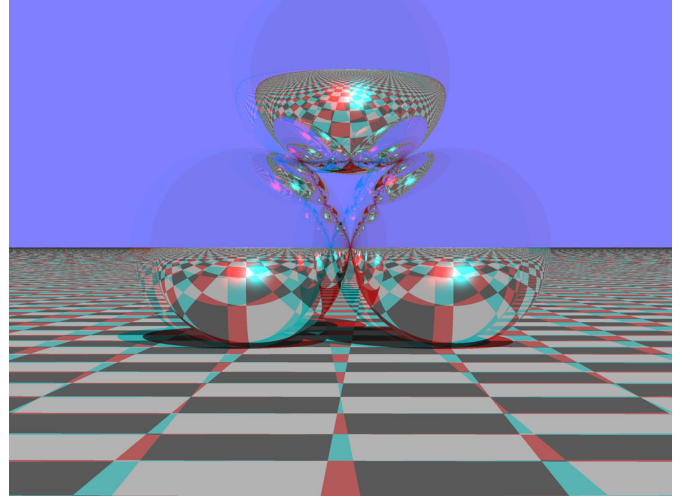


Figure 10: Example of a colour anaglyph image. In addition to the checkerboard floor, the scene contains three reflective spheres.

Figures 9 and 10 show examples of anaglyph images.

7. Random-dot autostereograms

TIM can render scenes as random-dot autostereograms [12]. Fig. 11 shows an example of such a random-dot autostereogram created by TIM.

Autostereograms rely on the fact that two different patterned surfaces can look identical from different viewing positions. In the case of standard autostereograms, the two different viewing positions are the positions of the observer's eyes, one of the surfaces is planar, and the pattern in that plane is the autostereogram of the other, three-dimensional, surface. For the observer to perceive the visual illusion of seeing the three-dimensional surface when viewing the autostereogram requires the patterns to be sufficiently detailed.

Placing dots on the two surfaces such that the patterns are consistent with each other is perhaps the simplest way to construct autostereograms. TIM uses the following algorithm. We call the plane of the autostereogram A , and the three-dimensional surface S (see Fig. 12).

1. Randomly pick a dot colour.
2. Randomly pick a position \mathbf{D} in the plane A and place a small dot of the picked dot colour there.
3. Find the point \mathbf{P} on S that lies on the same line of sight as \mathbf{D} , as seen from the position of the left eye, \mathbf{L} . For the surfaces A and S to look identical from the position of the left eye, \mathbf{P} therefore has to have the same colour as \mathbf{D} , namely the dot colour picked above.
4. Find the point \mathbf{D}' on A that lies on the same line of sight as \mathbf{P} , as seen from the position of the *right* eye, \mathbf{R} . By the same argument as above, the colour of this point must also be the dot colour picked above. Therefore, place another small dot of the picked dot colour at \mathbf{D}' .

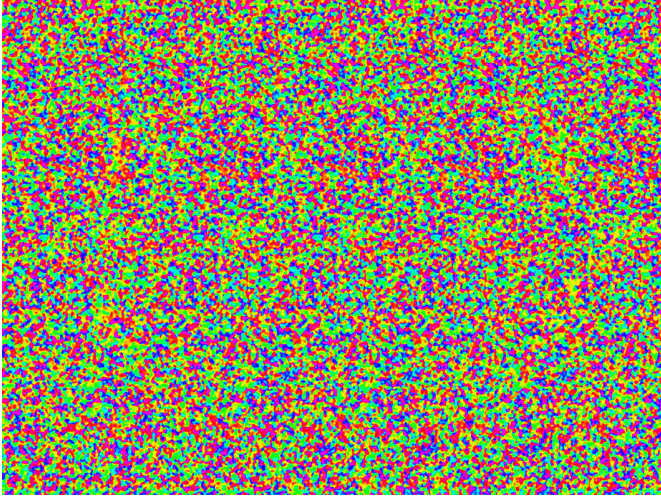


Figure 11: Random-dot autostereogram of TIM's default scene (Fig. 2). Note that the scene includes a plane behind Tim's head which can be seen in the background, making it easier to see the autostereogram.

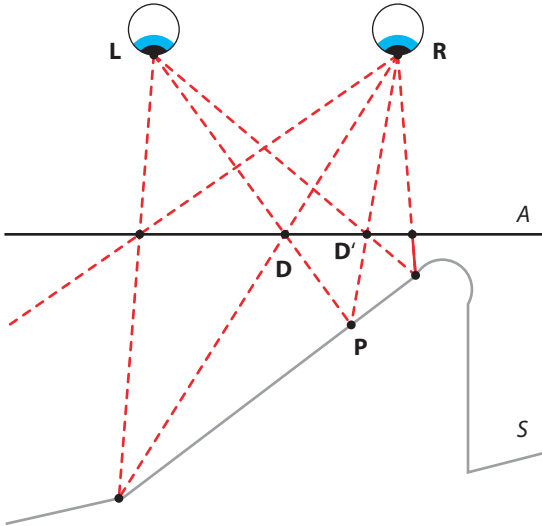


Figure 12: Principle of random-dot autostereograms. The eyes are located at the positions **L** and **R**; the autostereogram is in the plane **A**; **S** is the 3D surface represented by the autostereogram.

5. The previous two steps constructed, from the position **D** of one dot in the autostereogram, the position **D'** of another dot that has to have the same colour. Repeat the previous two steps to construct, from the position **D'** of this other dot, the position of yet another dot that has to have the same colour. Keep repeating until the position of the new dot lies outside the area of the autostereogram.
6. Steps 3 to 5 constructed, from the position **D** picked in step 2, the positions of further dots that have to be of the same colour. This was done by using the left eye's line of sight to construct a corresponding point on **S**, and then the right eye's line of sight to construct a point on **A** corresponding to this new point. Start again from the position picked in step 2, **D**, and repeat steps 3 to 5, but swapping the role of the left eye and the right eye. In other words, now use the *right* eye's line of sight to construct a point on **S** that corresponds to a dot position in **A**, and then use the *left* eye's line of sight to construct a point on **A** corresponding to this point on **S**, which is then the position of a new dot.

Two details in TIM's algorithm are perhaps worth noting.

1. The dots placed by TIM in the autostereogram are not single pixels but Gaussians. This means that their positions are not restricted to the pixel positions, so they can be placed "between pixels". This is advantageous as restricting dot positions to pixel positions restricts the possible separations between dots, and therefore the depths the dot patterns can represent, which results in "layered" random-dot stereograms.
2. Each dot TIM places in the autostereogram has a particular hue. To calculate the colour of a particular pixel in the autostereogram, TIM calculates the weighted average of the hues of all the dots that intersect at the pixel. Hue is the azimuthal angle ϕ of a colour's position in a suitable colour wheel (e.g. red = 0° , yellow = 60° , green = 120° , cyan = 180° , blue = 240° , purple = 300°). To form the weighted average of the hues due to all dots, TIM converts the hue of each dot into corresponding Cartesian coordinates (the coordinates corresponding to the hue ϕ_j of the j th dot, weighted by a factor w_j , are $x_j = w_j \cos(\phi_j)$ and $y_j = w_j \sin(\phi_j)$), adding up the x and y coordinates due to all hues, and calculating the azimuthal angle ϕ of the resulting (x, y) position. The weight of the hue of a particular dot is given by $w_j = \exp[-(r_j/\rho)^2]$, where r_j is the distance between the pixel and the centre of the j th dot, and ρ is the dot radius. This can be expressed in terms of complex numbers as

$$\phi = \arg \left\{ \sum_j \exp[-(r_j/\rho)^2] \exp(i\phi_j) \right\}. \quad (6)$$

Finally, the pixel colour is found using the standard Java method for converting a colour represented in terms of its hue, saturation and brightness (HSB) into its red, green and blue (RGB) components. Both saturation and brightness are set to their maximum value, 1.

8. Conclusions

TIM is a powerful raytracer with extensive capabilities, a number of them unique. We use TIM in our research on METATOYS: for disseminating our research over the internet by inviting playful experimentation with METATOYS through TIM's interactive version; and for conducting computer experiments with METATOYS ourselves.

Sometimes it is necessary to modify TIM's source code, which can be a daunting prospect. This paper is intended to help others and ourselves doing this. We hope it will entice other researchers to use TIM in their own work.

Acknowledgements

Thanks to The MathWorks and to NIST for making the source code of JAMA publicly available.

Appendix A. Implementation of ray tracing in TIM

TIM is written in Java [13], an object-orientated programming language [14]. As in all object-orientated languages, objects are instances of classes. A class has associated methods and data (variables), which take on specific values in specific instances of the class, i.e. the objects. A class can have subclasses, which inherit its methods and properties²; the former is called the latter's superclass. In Java there are also interfaces, which are collections of methods. If an object implements a specific interface, it has to implement all the interface's methods.

To trace a ray in TIM, a number of objects are interacting. The following are the most important ones.

Rays. Each ray is an object of class Ray, which has a starting position and a normalised direction, both represented as 3D Cartesian vectors.

Scene objects. Each scene object implements the interface SceneObject³; most (but, for complicated reasons, not all) are instances of a subclass of SceneObjectClass, which implements the SceneObject interface. Specifically, a scene object implements methods that calculate the intersection point between the scene object and any ray, and methods that calculate the colour of an intersection point when seen along the direction of a given ray and under specified illumination. There are particularly fundamental scene objects, each typically representing a simple geometrical object (such as a sphere) with specific surface properties (such as reflective), which are instances of the SceneObjectPrimitive subclass of SceneObjectClass. Another subclass of SceneObjectClass is SceneObjectContainer, which represents a collection of scene objects. There are also more complicated scene objects

that are neither SceneObjectPrimitives nor SceneObjectContainers, but ultimately every intersection with a SceneObject must be traceable to an intersection with a SceneObjectPrimitive.

Surface properties. Perhaps the simplest surface property is a colour that is independent of any light sources (effectively a coloured glow), which is represented by the class SurfaceColourLightSourceIndependent. The colour of non-glowing surfaces depends on illumination, and this is represented by the class SurfaceColour, in which the surface has separate diffuse and specular colours. The diffuse component colours light due to Lambertian reflectance [15]; the specular components colours light that is specularly (or near-specularly) reflected [16].

There are also classes representing surfaces on which the ray does not end, but which change its direction. Examples include mirror surfaces (Reflective) and refractive-index interfaces (Refractive).

Finally, there is currently one class (Teleporting) which continues tracing a light ray with a changed direction and from a new starting position.

All surface properties implement the interface SurfaceProperty.

Light sources. Light sources are represented by instances of the class LightSource. In TIM, two types of light source are implemented: ambient light, represented by the class AmbientLight, and light sources with a specific 3D position which throw shadows and which can create highlights due to specular (or near-specular) reflection off surfaces of type SurfaceColour.

Ray-tracing software usually considers only those light rays that eventually enter the virtual camera. They do this by tracing light rays *backwards*, starting from each pixel in the camera's virtual detector chip. What they try to do is establish the colour such a light ray would have if it was travelling in the opposite direction, i.e. the colour of the reverse ray, which is the colour an observer would see in the direction of the light ray. More details can be found in Ref. [17].

Backwards tracing a specific light ray in TIM proceeds as follows:

1. The SceneObject representing the entire scene is asked to return the colour of the reverse ray.
2. The SceneObject finds the SceneObjectPrimitive the light ray intersects (if any), and asks its SurfaceProperty to return the colour of the reverse ray.
3. The SurfaceProperty either changes the light-ray direction and starts tracing again, or it determines the colour of the surface under illumination by the LightSource. If the SurfaceProperty is a SurfacePropertyContainer, then the colours due to all the surface properties are summed (by adding up the individual RGB components).
4. The LightSource returns a colour according to its shading model. If the LightSource is a

²Unless they prevent this specifically by explicitly overriding individual methods and/or properties.

³Note that every Java class is a subclass of the Object class.

LightSourceContainer, then it asks each of the LightSources it contains to return a colour and then sums these colours.

We discuss the steps in some more detail below.

In TIM, tracing an individual ray backwards is initiated by asking the SceneObjectContainer containing all scene objects to return the colour an observer at the ray's starting point would see in the ray's direction. This is done by calling the SceneObjectContainer's getColour method (which is defined in the SceneObjectContainer's superclass SceneObjectClass). The SceneObjectContainer then establishes which one (if any) of the scene objects it contains the ray would intersect with first. If the ray intersects none of the scene objects, then the getColour method returns the colour black. If the ray intersects one of the scene objects, then the method establishes which SceneObjectPrimitive was hit and calls this SceneObjectPrimitive's SurfaceProperty to establish the colour.

Each SurfaceProperty implements a getColour method, which returns the colour an observer at the ray's starting point would see in the ray's direction. In the simplest case, implemented in the class SurfaceColourLightSourceIndependent (surface type "Coloured (glowing)"), the colour stored within the specific instance of the class is returned, irrespective of the illumination and the rest of the scene. Illumination-dependent colour is handled by the SurfaceColour class, which calls the light source's getColour method to establish the appearance of the surface under illumination by the light source. There are also surface properties, for example Reflective, whose getColour method returns the colour resulting from tracing through the scene a new ray that starts from the intersection point and travels in a new direction (given, in the case of the Reflective class, by the law of reflection). Finally, it is also possible to have surface properties whose getColour method returns the colour resulting from tracing a new ray that starts at a point different from the intersection point; the Teleporting class is an example of such a class. In the latter two cases of surface properties that continue tracing rays through the optical system, the colour may be slightly darkened to represent a reflection or transmission coefficient of less than one.

A light source's getColour method calculates the colour in which an observer would see a specific surface colour at a particular point on a surface if it was illuminated only by this light source. TIM models two different types of light source: ambient light (class AmbientLight), which illuminates all objects in all directions with a given RGB colour; and PhongLightSource, which implements the Phong shading model [16]. The latter corresponds roughly to a small, coloured light bulb: it has a specific position, which is used to determine whether the surface is in another scene object's shadow and whether near-specular reflection occurs, which leads to highlights; and it has a specific colour. There is also a class (LightSourceContainer) that models the effect of combinations of different light sources.

It should be clear from the above discussion that

backward tracing of a ray ends when a light ray has intersected a surface with a surface of class SurfaceColourLightSourceIndependent, or when it hits a surface of class SurfaceColour (in which case the light source performs the final calculation of colour). Sometimes it can happen that rays get "trapped", for example between mirror surfaces such as those in canonical optical resonators [18]. In such cases, a reflection (or transmission) coefficient < 1 ensures exponential fall-off of the intensity with the number of bounces, so such light rays become dark. TIM limits the number of bounces it models, and when the limit is reached returns the colour black. This is controlled by a variable called traceLevel, which gets passed between the objects. The backwards tracing process starts with traceLevel taking a value of typically 100; whenever a surface property initiates tracing of a new ray, it does so with a value of traceLevel that is reduced by 1. When the value 0 is reached, the colour black is returned.

Appendix B. Source-code structure

TIM is divided into a hierarchical package structure. We describe here the main branch of this structure, namely the optics package and the packages it contains. There are three additional packages:

1. JAMA is a matrix package in the public domain [19]. TIM uses it; for convenience, the unmodified third-party source code is distributed with TIM's source code.
2. math is a package that defines additional mathematical functionality as required, including classes dealing with complex numbers, 2D vectors, and 3D vectors.
3. test contains only the class Test, which can be executed as a Java application for the purposes of testing any parts of the code.

Appendix B.1. The optics package

The optics package collects together optics-related code. At the top level, it contains the Constants class, a collection of optics-related constants such as a few common refractive indices; and DoubleColour, which is used internally to represent colours.

The only sub-package within the optics package that is distributed with TIM's source code is optics.raytrace. It contains both the ray tracer code as well as associated mathematical, graphical and user interface code, organised in the form of a number of sub-packages. At the top level, it contains NonInteractiveTIM, a template for a class that can be run as a Java application that uses TIM (see section Appendix C); TIMApplet, the applet class that is called when the interactive version of TIM is run as an applet; TIMJavaApplication, which allows the interactive version of TIM to be run as a Java application (with slightly increased functionality, specifically the ability to save images as .bmp files); and TIMInteractiveBits, a class that defines the interactive parts of TIM, which are called by both the TIMApplet and TIMJavaApplication classes.

The package `optics.raytrace.core` contains a number of the core ray-tracing classes and interfaces. A number of these, and their interactions, are discussed in Appendix A. The ray-tracing core classes and interfaces include those defining the structure of cameras (`Camera`), light sources (`LightSource`; see `optics.raytrace.lights` for implementations), scene objects (`SceneObject`; implementations are in package `optics.raytrace.sceneObjects`), surface properties (`SurfaceProperty`; implementations in `optics.raytrace-surfaces`), rays (`Ray` and `RayWithTrajectory`), and intersections between rays and objects (`RaySceneObjectIntersection`). The `Studio` class defines a collection of everything required to calculate a photo, namely information about the scene (in a `SceneObject` object), lighting (a `LightSource` object), and camera (in the form of a `Camera` object; implementations of cameras are in the package `optics.raytrace.cameras`). A number of interfaces outline mechanisms for the parametrisation of object surfaces:

- `ParametrisedObject` defines coordinates assigned to each point on a surface;
- `One2OneParametrisedObject` extends `ParametrisedObject` by asking for the reverse of the coordinate assignment, i.e. for a method that returns the point on a surface associated with a set of coordinate values.

Other classes included in the package are `CameraClass`, which implements a number of methods common to currently all cameras, and which is a superclass to currently all cameras; `SceneObjectClass`, which similarly implements a number of methods common to most scene objects, and is a superclass of many scene-object classes; `SceneObjectPrimitive`, which describes simple geometric objects, such as spheres and planes; `CCD` and `CentredCCD`, which represent the light-detecting element in cameras; `Transformation`, which defines the structure of geometrical transformations (such as translation or rotation) which can be applied to scene objects (for implementations see package `optics.raytrace.sceneObjects.transformations`); and `DoubleColourCamera` and `DoubleColourCCD`, which define the structure of higher-quality cameras and their light-detecting elements.

The useful `optics.raytrace.demo` package contains a number of classes which can be run as Java applications and which demonstrate the use and effect of different features, for example `LightsDemo`, which was used to create Fig. B.13. It is worth studying the examples contained in this package to understand how to access TIM's functionality.

The package `optics.raytrace.exceptions` defines classes that signal exceptional circumstances that occurred during rendering, for example a ray becoming evanescent (`EvanescentException`).

Appendix B.2. *optics.raytrace.cameras*

The `optics.raytrace.cameras` package contains implementations of the `Camera` interface. These handle the

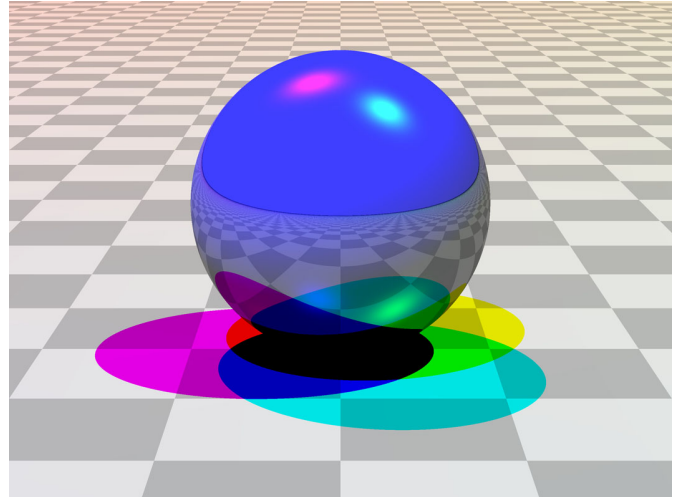


Figure B.13: Effect of combination of different light sources. Here, a shiny blue sphere is illuminated by a combination of three Phong light sources, one red, one green, one blue, placed in different directions high above the sphere. The Phong light sources produce differently-coloured highlights at the top of the sphere and differently-coloured shadows (the colours are due to subtractive colour mixing) on the floor. Where the shadows from all Phong light sources overlap, the scene is completely black in the absence of an ambient light source. The image was rendered using the `LightsDemo` class in the `optics.raytrace.demo` package (Appendix B).

mechanisms of generating the rays that are then traced through the scene, and of turning creating corresponding images. Implemented camera classes include `PinholeCamera`, which is the simplest type of camera that takes pictures in which everything is in focus; `ApertureCamera`, a camera with a circular aperture that can focus on any transverse plane; `AnyFocusSurfaceCamera`, a camera with a circular aperture that can focus very general surfaces (section 3); `OrthographicCamera`, which produces orthographic projections into a plane; `AnaglyphCamera`, which can produce either red/blue or colour anaglyph images that can be viewed with standard red/cyan anaglyph glasses (section 6); and `AutostereogramCamera`, which can create random-dot autostereograms of the scene (section 7).

Appendix B.3. *optics.raytrace.lights*

The `optics.raytrace.lights` package includes implementations of light sources. These include `AmbientLight`, which represents a coloured ambient light; `PhongLightSource`, which realizes the Phong shading model [16], which is roughly equivalent to a slightly fuzzy point light source; and `LightSourceContainer`, which allows light sources to be combined. Fig. B.13 demonstrates effects due to different light sources.

Appendix B.4. *optics.raytrace.sceneObjects*

The package `optics.raytrace.sceneObjects` contains implementations of scene objects. It contains classes describing different types of scene objects:

Simple geometrical shapes. Classes that describe simple geometrical shapes are implementations of the `SceneObjectPrimitive` class (Appendix D.1). Examples include spheres (`Sphere`), planes (`Plane`), parallelograms (`CentredParallelogram`), discs (`Disc`), and cylinders (`Cylinder`).

Combinations of other scene objects. A number of classes describe compound objects (Appendix D.2). Examples include `Arrow`, `Cylinder`, and `Eye`.

Shapes with parametrised surfaces. Parametrisation, described by implementations of the interfaces `ParametrisedObject` and `One2OneParametrisedObject`, assigns coordinates to points on the surface of a geometrical shape (Appendix D.3). Examples of classes that define parametrised geometrical shapes include `ParametrisedSphere`, `ParametrisedPlane`, and `ParametrisedCentredParallelogram`. There are also classes that allow the range of the coordinates that describes the surface of the geometrical shape to be varied. For example, in the `ScaledParametrisedSphere` class, the range of the polar angle θ can be set to range from an arbitrary value θ_{\min} to another arbitrary value θ_{\max} . Other examples include `ScaledParametrisedDisc` and `ScaledParametrisedCentredParallelogram`.

The `optics.raytrace.sceneObjects.solidGeometry` package is a collection of classes useful for combining scene objects. In the simplest case, scene objects are grouped in a hierarchical way (`SceneObjectContainer`). More elaborate combinations of scene objects include intersections (`SceneObjectIntersection`), unions (`SceneObjectUnion`), and inverse (`SceneObjectInverse`).

One of the capabilities required of any scene object is the ability to create a transformed copy of itself. The structure of the transformation is defined by the `Transformation` class (in `optics.raytrace.core`). The package `optics.raytrace.sceneObjects.transformations` contains classes that describe specific types of transformation, i.e. subclasses of `Transformation`. Examples include `Translation`, `RotationAroundXAxis`, `RotationAroundYAxis`, `RotationAroundZAxis`, and the more general `LinearTransformation`.

Appendix B.5. *optics.raytrace.surfaces*

The `optics.raytrace.surfaces` package contains implementations of the `SurfaceProperty` interface (in `optics.raytrace.core`). These include

- the classes representing coloured surfaces, `SurfaceColourLightSourceIndependent` and `SurfaceColour`;
- a class representing a transparent surface (`Transparent`);
- the `Reflective` class which represents specularly reflective surfaces;

- the `Refractive` class which represents refraction at the interface between media with different refractive indices according to Snell's law;
- a class that facilitates the implementation of classes that represent surfaces that change direction according to generalised laws of refraction (`Metarefractive`);
- a number of subclasses of `Metarefractive` representing surfaces that invert one of the ray-direction components tangential to the surface [20] (`RayFlipping`), rotate the ray direction around the local surface normal [5] (`RayRotating`), and refract like — formally [4] — the interface between media with a *complex* refractive-index ratio would (`RefractiveComplex`);
- classes representing surfaces that have combinations of other surface properties (`SurfacePropertyContainer` and `SurfacePropertyContainerWeighted`);
- a `SemiTransparent` class, which combines an arbitrary surface property with the `Transparent` surface property;
- a class representing a surface whose inside has different properties from its outside (`TwoSidedSurface`);
- classes representing surfaces with completely different surface properties at different points on the surface (`SurfaceTiling`, `EitherOrSurface`, `PictureSurface`, `PictureSurfaceDiffuse`, `PictureSurfaceSpecular`);
- a class representing the hologram of a thin lens (`ThinLensHologram`), which changes light-ray direction dependent on the point where it is being intersected;
- and a class representing a surface a light ray enters and then continues, from a corresponding position and with a corresponding direction, from the surface of a different scene object (`Teleporting` — see section 4; we use this to model geometric optical transformations [8]).

Appendix E.1 discusses how to add new surface properties to TIM.

The `optics.raytrace.surfaces.metarefraction` package contains description of the direction change performed by surfaces of class `Metarefractive` (see `optics.raytrace.surfaces`). The format of these descriptions is defined by the abstract `Metarefraction` class. The `ComplexMetarefraction` class is a subclass of `Metarefraction`, again abstract, which formulates the direction change in terms of the projections of the incoming and outgoing light rays into an Argand plane tangential to the surface at the intersection point and with its origin there (section 2) [6]. This is a generalisation of the formal description in terms of multiplication with a complex number of rotation of the light-ray direction around the local surface normal [4]. All other classes in this package are non-abstract subclasses of `ComplexMetarefraction`. They include classes that describe surfaces that change light-ray direction described by complex

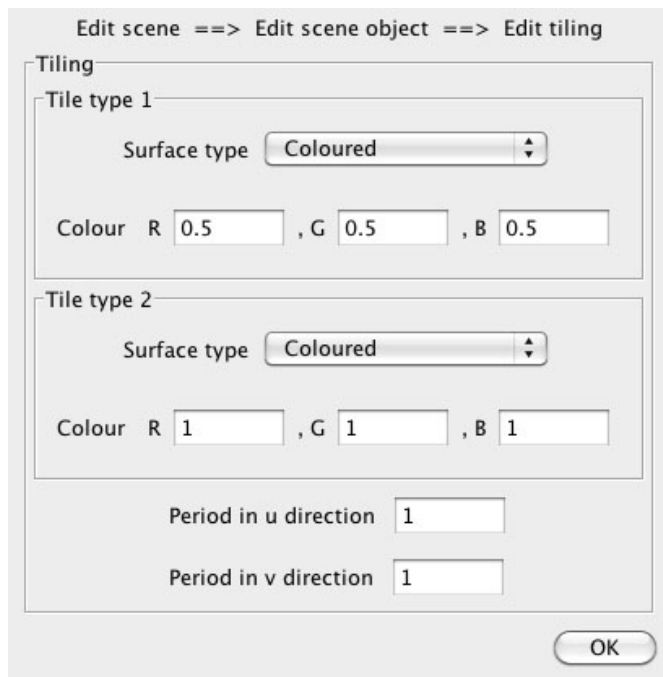


Figure B.14: Editing one object in a hierarchical structure using the `IPanel` class. The top line displays a breadcrumb trail, giving an idea of the place in the hierarchy.

multiplication (`ComplexMetarefractionMultiplication`), complex addition (`ComplexMetarefractionAddition`), complex conjugation (`ComplexMetarefractionCC`), and complex exponentiation (`ComplexMetarefractionExp`).

Appendix B.6. *optics.raytrace.GUI*

The `optics.raytrace.GUI` package is a collection of packages that together constitute TIM's graphical user interface (GUI). All but one of the sub-packages of `optics.raytrace.GUI` contain classes that handle user interaction in TIM's interactive version.

The exception that contains the classes that handle user interaction in the non-interactive version of TIM is the package `optics.raytrace.GUI.nonInteractive`. It contains two classes: `PhotoCanvas`, which creates a panel (screen element) in which `NonInteractiveTIM` and classes derived from it display the rendered image; and `PhotoFrame`, which opens a window containing a `PhotoCanvas`.

The package `optics.raytrace.GUI.core` contains GUI core classes. The `RaytraceWorker` class handles rendering in a background thread, so that the GUI can continue to react to user interaction. The class `IPanel` defines a screen element intended to allow browsing and editing hierarchical networks of objects on a relatively small screen area, so that it can be integrated into an internet page without necessarily dominating the page. At any one time, the panel corresponding to one object in this network is displayed. `IPanel` handles a stack of these panels, displaying only the top one, and a breadcrumb trail of the other panels in the stack (Fig. B.14). The mechanism by which this panel is supplied by the object, and how control is

handed over when another object's panel gets displayed in the `IPanel`, is defined in the `IPanelComponent` interface, which all objects in TIM which are editable interactively implement. `EditableCamera` is an interface that defines the functionality of an editable camera.

The package `optics.raytrace.GUI.lowLevel` is a collection of low-level classes related to the GUI. The class `GUIFrame` represents the window TIM's interactive version opens when run as a Java application. `GUIPanel` is the top-level screen element that contains the entire GUI, i.e. the tabs showing views of the rendered scene from different view points, any component being edited, and all buttons (see Fig. 2). `BufferedImageCanvas` is the panel displaying the rendered image. `RaytracingImageCanvas` extends `BufferedImageCanvas` by adding some interactive functionality, including the ability to identify the scene object and coordinates of the point on the object the mouse pointer hovers over and clicks on, and the ability to edit that scene object by double-clicking on it. Most classes provide panels for editing small chunks of data, including integers (`IntPanel` and `LabelledIntPanel`, an `IntPanel` displayed next to a brief verbal description), individual double-precision real numbers (`DoublePanel` and `LabelledDoublePanel`), pairs of double-precision real numbers (`TwoNumbersPanel`), complex numbers (`ComplexPanel` and `LabelledComplexPanel`), 2D vectors (`Vector2DPanel` and `LabelledVector2DPanel`), 3D vectors (`Vector3DPanel` and `LabelledVector3DPanel`), and the limiting values of a range of real numbers (`LabelledMinMaxPanel`). The class `SceneObjectPrimitivesComboBox` describes a panel that allows any `SceneObjectPrimitive` in the scene to be selected, which is used in the panel editing surface properties for selecting a target object for a surface of type `Teleporting`. There are a number of classes for editing camera-specific parameters such as aperture size (`ApertureSizeComboBox` and `LabelledApertureSizeComboBox`); quality (`QualityComboBox` and `LabelledQualityComboBox`), which can be applied to blurring and to anti-aliasing; and all parameters related to camera blur collected in one panel (`BlurPanel`). The class `ButtonsPanel` describes a panel containing one or more buttons. The interface `StatusIndicator` outlines the structure of an object that can display brief status messages. Finally, `GUIBitsAndBobs` is a small collection of miscellaneous methods commonly used by the GUI.

The package `optics.raytrace.GUI.cameras` contains classes that describe various editable camera classes. The classes `EditableAnyForucSurfaceCamera`, `EditableApertureCamera`, `EditableAnaglyphCamera`, `EditableAutostereogramCamera` and `EditableOrthographicCamera` respectively extend the `AnyFocusSurfaceCamera`, `ApertureCamera`, `AnaglyphCamera`, `AutostereogramCamera` and `OrthographicCamera` classes in this way. The classes `EditableOrthographicCameraSide` and `EditableOrthographicCameraTop` make special cases of the `OrthographicCamera` class editable; these respectively correspond to TIM's "Side view" and "Top view" tabs.

The package `optics.raytrace.GUI.sceneObjects` contains the classes describing all editable scene objects. Most of these are simply editable versions of classes in `optics.raytrace.sceneObjects`, including `EditableArrow`, `EditableParametrisedCone`, `EditableParametrisedCylinder`, `EditableParametrisedPlane`, `EditableRayTrajectory`, `EditableRayTrajectoryCone`, `EditableScaledParametrisedDisc`, `EditableScaledParametrisedCentredParallelogram`, and `EditableScaledParametrisedSphere`. The class `EditableSceneObjectCollection` allows editing of groups of scene objects, which can be combined a number of ways respectively handled by the `SceneObjectContainer`, `SceneObjectIntersection`, and `SceneObjectUnion` classes in the `optics.raytrace.sceneObjects.solidGeometry` package. A few of the classes defined in `optics.raytrace.GUI.sceneObjects` exist only in editable form. Examples include `EditableCylinderFrame`, `EditableCylinderLattice`, `EditableLens`, `EditableObjectCoordinateSystem`, `EditablePolarToCartesianConverter`, and `EditableTelescope`.

The `optics.raytrace.GUI.surfaces` package contains classes that enable selecting a class of surface property and editing class-specific parameters. The class that provides a panel for doing all of this is `SurfacePropertyPanel` (see Fig. E.15, appendix Appendix E.3). The remaining classes in this package allow editing of class-specific parameters. `EditableSurfaceTiling` and `EditableTwoSidedSurface` are editable subclasses of `SurfaceTiling` and `TwoSidedSurface`, respectively. `TeleportingTargetsComboBox` is a subclass of `SceneObjectPrimitivesComboBox` (in the `optics.raytrace.GUI.lowLevel` package) that allows a suitable scene object contained in the scene to be selected. In its labelled form (`LabelledTeleportingTargetsComboBox`) this is used in the `SurfacePropertyPanel` class to select a target object for the `Teleporting` surface property.

Last, and least, the package `optics.raytrace.GUI.sceneObjects.transformations` contains only the class `EditableLinearTransformation`, which will eventually be able to edit linear scene-object transformations (see package `optics.raytrace.sceneObjects.transformations`) and become part of TIM's interactive version.

Appendix C. The default non-interactive TIM

TIM's source code comes with a class called `NonInteractiveTIM`, which can be compiled and run as a Java application⁴ that defines a studio (scene, camera and

lights); renders the scene under the conditions defined in the studio; and displays the rendered image on the screen and saves it as a .BMP file. This class can serve as an example and template for using TIM's source code.

The class `NonInteractiveTIM` provides three methods:

1. `getFilename` returns the filename under which the rendered image is saved;
2. `createStudio` defines and returns the studio, i.e. a scene, a camera, and lights;
3. `main` calls the `createStudio` method, renders the image, displays it on the screen, and saves it as a .BMP file; automatically gets called when the `NonInteractiveTIM` class is run as a Java application.

Modifying these methods⁵ allows TIM to perform tasks that cannot currently be achieved in the interactive version; below are a few examples.

First we discuss modifying the `createStudio` method, which changes one or more of scene, camera and lights. This allows the programmer to do a number of things that are not currently possible in the interactive version, including

- setting parameter values by typing in formulas;
- using Java's built-in flow control (such as `for` or `while` loops) and formulas to create a number of objects systematically;
- accessing classes of scene object not currently supported in the interactive version (e.g. `MaskedObject`);
- transforming objects, e.g. rotate, move, or scale them (these transformations are defined in the `optics.raytrace.sceneObjects.transformations` package);
- accessing additional classes of surface properties (an example is `PictureSurface`, which maps a picture loaded from a file onto a surface; other examples include several types of `MetarefractiveSurface`);
- giving surfaces *combinations* of surface properties (which can be achieved by using the `SurfacePropertyContainer` or `SurfacePropertyContainerWeighted` classes);
- changing the lights.

Changing the behaviour of the resulting Java application altogether can be achieved by altering the main class. Simple examples include stopping the Java application from saving the image (by removing the relevant line in the code, or simply by commenting it out). By running a loop in which parameters change values, for example the position of a scene object or the camera, and by saving the rendered images with a suitable filename (in the simplest case ending in a number), the saved images can later be combined into a movie by other software.

⁴In Eclipse [21], simply bring up the source-code file `optics.raytrace.RayTraceJavaApplication.java` and select `Run > Run`.

⁵Note that overriding these methods has no effect as all are declared `static`.

Appendix D. Adding a new scene-object class

Sometimes it is necessary to add a new class of scene object to TIM. We can distinguish the following cases, which we treat in more detail in the following sections:

1. It is desirable to add a class representing a geometrical shape not yet represented in TIM, for example an ellipsoid or a torus.
2. It is desirable to define a class representing combinations of geometrical shapes already represented in TIM. One reason for doing so could be to automate the placing and adding to the scene of the constituent scene objects, for example the cone and the cylinder that form an arrow represented by the Arrow class. The individual scene objects can also automatically be combined using the solid-geometry classes defined in the `optics.raytrace.sceneObjects.solidGeometry` package; this is how a lens is created by the EditableLens class.
3. It is desirable to (re)parametrise the surface of an existing geometrical shape.
4. It is desirable to add a scene-object class that is represented in TIM for non-interactive use to the interactive version.

Appendix D.1. Adding a class representing a geometrical shape

A geometrical shape is represented by a (non-abstract) subclass of the (abstract) `SceneObjectPrimitive` class. This is itself a subclass of `SceneObjectClass`, an abstract class that implements some common methods required by the `SceneObject` interface such as keeping a copy of the studio, the parent object, and the description, and providing implementations of methods such as `getColour` that reduce the task to other methods that remain to be implemented, such as finding the intersection between a ray and the object. It is instructive to study the implementation of such a class, for example `Sphere`.

Any new subclass of the `SceneObjectPrimitive` class needs to implement methods for dealing with the geometry of ray tracing, namely finding the intersection between a ray and the shape (`getClosestRayIntersection`); calculating the normalised surface normal at any possible intersection point (`getNormalisedSurfaceNormal`); and determining whether a position is inside the object or not (`insideObject`). It also needs to implement methods for an instance of the class to make an identical copy of itself (`clone`), or a copy that is geometrically transformed, for example shifted, rotated, or scaled (`transform`).

Note that the surface normal points in the direction of the shape's outside. In many cases, it is obvious which side of the surface is on the inside and which one is on the outside, for example in the case of a sphere. However, in other cases, for example in the case of a plane, it is not at all obvious. In such cases, the direction of the surface normal *defines* an inside and an outside. It is important to be able to distinguish inside from outside as many types of surface property, for example refraction, distinguish between rays that arrive from the inside from those arriving from the outside.

Appendix D.2. Defining a new scene-object class in terms of existing scene objects

Sometimes it is desirable to define a new class of scene objects that consists of a number of scene objects of existing classes. Examples include the Arrow class, which represents a combination of a cone (the arrow's tip) and a cylinder (the shaft), and the EditableLens class, which represents a convex-convex lens.

The easiest way to implement such a class is to extend the class that represents the appropriate combination of scene objects. For example, an arrow, which is simply a collection of a cone and a cylinder, can be realised by extending the class representing simple scene-object collections, `SceneObjectContainer`; a convex-convex lens, which is the intersection of two spheres, can be realised by extending the `SceneObjectIntersection` class. All the class then needs to do (usually in the constructor) is to add the appropriate scene objects to the array of objects in the class, using the `addSceneObject` method.

Appendix D.3. (Re)parametrising the surface of an existing shape

A number of surface properties require a surface that has a two-dimensional coordinate system associated with it (section 4, especially Fig. 6). An example of such a surface property is `SurfaceTiling`, which covers the surface in a chequer-board pattern in the surface's coordinate system. A surface that can calculate a pair of coordinates for any point on the surface is represented by the `ParametrisedObject` interface; it needs to implement methods for returning coordinates of an arbitrary point on the surface (`getSurfaceCoordinates`) and for returning the corresponding coordinate names, for example `theta` and `phi` in the case of a sphere's polar coordinate system (`getSurfaceCoordinateNames`). A surface that can also identify a point on the surface for arbitrary coordinates is represented by the `One2OneParametrisedObject` interface. Such a surface needs to implement the methods required by the `ParametrisedObject` interface, and additionally the `getPointForSurfaceCoordinates` method. If the mapping between surface points and coordinates is not one-to-one then the behaviour of this method is undefined.

Note that there are scene objects, most notably compound scene objects, which can implement the `ParametrisedObject` interface so that patterned surfaces can be applied to each constituent scene object, provided it implements `ParametrisedObject`, but which should not implement the `One2OneParametrisedObject` interface even if all constituent scene objects do as there might be points on different constituent objects that correspond to the same combination of coordinate values, and so the mapping between surface points and coordinates is not one-to-one.

A scene object that implements the `ParametrisedObject` interface also defines directions on the surface, through the `getSurfaceCoordinateAxes` call-back method. This allows the implementation of anisotropic surface properties, which require a preferred direction to be defined. Fig. 6 shows these vectors for two points on the surface of an object of class

ParametrisedSphere. Their primary purpose is to define directions on the surface, which is why they many methods that use these vectors normalise them. An example is the surface-property class *Metarefractive*, which allows surfaces to deflect light rays in very general ways. (This can be seen as a generalisation of refraction at the interface between media with different refractive indices, or “metarefraction” [1].) As many light-ray-direction changes are internally handled through the *Metarefractive* class (see Appendix E.2), it is often important that surfaces implement the *ParametrisedObject* interface. Those light-ray-direction changes include Snell’s-law refraction (surface property *Refractive*); ray flipping, which changes the sign of one light-ray-direction component [20] (surface property *RayFlipping*); and ray rotation (Fig. 2), which rotates the light-ray direction by an arbitrary (but fixed) angle around the local surface normal [5] (surface property *RayRotating*).

Creating a class that parametrises the geometrical shape described in an existing class can be achieved by creating a subclass of the existing class. This new subclass needs to implement the methods of the relevant interfaces, *ParametrisedObject* or *One2OneParametrisedObject*. If the existing class is already parametrised and the new subclass overrides its methods related to parametrisation, then the new class re-parametrises the geometrical shape described in the existing class.

Appendix D.4. Adding an existing scene-object class to interactive TIM

All scene-object classes available in the interactive version of TIM need to be fully parametrised (see previous section), i.e. they need to implement the *One2OneParametrisedObject* interface.

For a class to be editable through the mechanism built into the interactive version of TIM, it needs to implement the *IPanelComponent* interface. To add an existing scene-object class to the interactive version of TIM, it is easiest to create a subclass that implements the following methods required by *IPanelComponent*:

1. *createEditPanel* prompts the creation of the edit panel (of class *JPanel*), an area of screen that allows interactive editing of any of the scene object’s parameters;
2. *discardEditPanel* signals to the *IPanelComponent* that the edit panel is no longer required, and any resources associated with it can be freed up;
3. *getEditPanel* returns the edit panel;
4. *setValuesInEditPanel* sets all the sub-panels in the edit panel to reflect the current values of the scene object’s parameters;
5. *acceptValuesInEditPanel* sets the scene object’s parameters to the edited values in the edit panel;
6. *backToFront* gets invoked after editing of a different *IPanelComponent*’s edit panel has been completed (for example, the surface property of the current scene object could have been edited), and this one’s edit panel is being edited again.

The new class can now be part of the scene, and it can be edited, but it is so far not possible to create a new instance (unless another one is being duplicated). Creating a new instance of a scene object in interactive TIM happens by adding a new scene object to a collection of scene objects. This collection can be “The scene”, which is the top-level collection of scene objects, but it can also be a collection that is part of “The scene” (or of other collections in the hierarchy). Editing a collection is handled by the *EditableSceneObjectCollection* class. We will discuss the required steps using the example of the *EditableTelescope* class.

1. Add a string that describes an object of this class (e.g. `String OBJECT_TELESCOPE = "Telescope";`).
2. Add this string to the array of strings that are to be the menu items in the combo box responsible for initiating the creation of a new instance of a scene object (e.g. `String[] componentStrings = { [...], OBJECT_TELESCOPE, [...]};`)
3. In the *actionPerformed* method of the internal class *SceneObjectControlPanel*, which gets called whenever the user interacts with the combo box that initiates the creation of a new scene object, add a case that gets invoked when the user has selected the new object type in the new-object combo box. In the case of our example, it has the following form:

```
else if(newElementName.equals(
    OBJECT_TELESCOPE)
)
    iPanelComponent = new EditableTelescope(
        "Telescope", // description
        // default centre
        new Vector3D(0, 0, 10),
        // default ocular normal
        new Vector3D(0, 0, -1),
        1, // default magnification
        1, // default radius of aperture
        // parent in hierarchy
        EditableSceneObjectCollection.this,
        // scene, lights and camera
        getStudio()
    );
```

Appendix E. Adding a surface-property class

One of the main reasons why we wrote TIM was to be able to fully control the effect of surfaces on light rays, which has proved tremendously useful not only for our research on METATOYS but also for our work on optical orbital angular momentum. This section outlines how new surface properties can be added to TIM.

Appendix E.1. Adding a general surface property

TIM establishes the effect of a surface on any specific light ray by asking the *getColour* method of the *SurfaceProperty* object representing the surface to return the colour of the reverse ray, i.e. the light ray travelling in the opposite direction.

The role of surfaces in establishing this colour is outlined in more detail in Appendix A.

A new surface property can be created by implementing the `SurfaceProperty` interface directly. The class representing the new surface property must implement the `getColour` method. Precisely how it calculates the reverse ray's colour varies greatly between different surface-property classes. The `getColour` method has access to all the information being passed to it as arguments:

- the `Ray` object describing the incident light ray contains its direction and starting position;
- the `RaySceneObjectIntersection` object describing the intersection between the surface and the light ray contains the position of the intersection point and the primitive scene object being intersected;
- the entire scene information is passed in the form of a `SceneObject` object;
- information about lights is passed in the form of a `LightSource` object.

Additionally, the trace level is passed as an argument. The class might also store additional information and make it available to be used within `getColour`.

It might be instructive to study the code of a few classes that implement the `SurfaceProperty` interface directly. The following two classes are perhaps particularly instructive.

- The class `SurfaceTiling` is an example of a spatially varying surface property that calculates the local coordinates in the intersected primitive scene object's coordinate system. This is done by the code fragment

```
(
    (ParametrisedObject)(i.o)
).getSurfaceCoordinates(i.p),
```

where `i` is the `RaySceneObjectIntersection` object describing the intersection point; `i.o` is the primitive scene object being intersected (which has to implement the `ParametrisedObject` interface here; if it does not, an exception is thrown); and `i.p` is a 3D vector describing the position of the intersection point.

- The class `Reflective` is an example of a surface property that requires further ray tracing. Its `getColour` method illustrates, amongst other things, checking that the trace level is greater than zero and continuing backwards ray tracing with a new direction. The latter is achieved by the following code segment:

```
return scene.getColourAvoidingOrigin(
    ray.getBranchRay(i.p, newRayDirection),
    // the primitive scene object being
    // intersected
    i.o,
```

```
1,    // the light source(s)
scene, // the entire scene
traceLevel-1
).multiply(reflectionCoefficient);
```

Note that creating the continuation of the ray using the original ray's `getBranchRay` method ensures the ray trajectory is recorded correctly (see section 5); that the branch ray is being launched with a trace level reduced by 1; and that the intensity of the branch ray is multiplied by a reflection coefficient stored in the `Reflective` object.

Appendix E.2. Adding a surface property that affects only light-ray direction

As surface properties that change light-ray direction in different ways are of particular interest to us for the purposes of our METATOYS research, we have put in place a surface-property class that aims to facilitate the creation of new light-ray-direction-changing surface properties. The direction change itself is described by a subclass of the abstract `Metarefraction` class, which must provide methods that calculate the new light-ray direction from the old light-ray direction for the two cases of light travelling inwards, i.e. arriving from the surface's outside, or outwards, i.e. arriving from the inside. (Inside and outside are defined by the direction of the normalised surface normal, which can be obtained with the `SceneObjectPrimitive` class's `getNormalisedSurfaceNormal` method, and which is defined to point outwards.) These methods are the call-back methods `refractInwards` and `refractOutwards`. The surface-property class that represents this light-ray-direction change is of class `Metarefractive`. The description of the direction change (i.e. the implementation of the `Metarefraction` class) is passed as an argument to the constructor of the `Metarefractive` surface property; in other words, an object of class `Metarefractive` is always created around a specific light-ray-direction change.

In TIM, most light-ray-direction changes at surfaces are described in terms of their projection into an Argand plane tangential to the surface at the intersection point and with its origin placed there (section 2) [6]. A complex number that corresponds to the projection of the incoming light ray is mapped to another complex number, which corresponds to the projection of the outgoing light ray. The component in the direction of the surface normal is calculated so that the length of the direction vector remains unchanged. The mapping of the light-ray projection then defines the direction change; for example, refraction according to Snell's law is described by multiplication with a real number, and rotation through an angle α around the surface normal is described by multiplication with a complex number of the form $\exp(i\alpha)$ [4]. Direction changes can be described in this way by the `ComplexMetarefraction` class, which is a subclass of `Metarefraction`.

Examples of surface properties which have been implemented by extending the `Metarefractive` surface-property class, and which use extensions of the

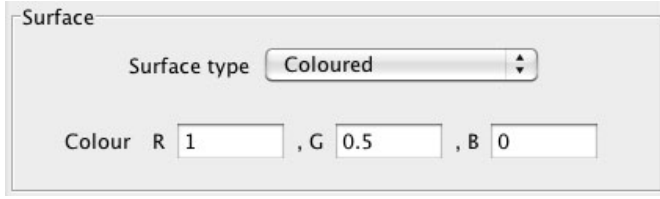


Figure E.15: The panel for editing the surface property of an orange object. The top half allows selecting a surface-property class; the lower half allows editing of parameters specific to the selected class, in this case the red (R), green (G), and blue (B) values of the surface colour. The panel is created by the class `SurfacePropertyPanel`.

`ComplexMetarefraction` class to describe the light-ray-direction change they represent, include `RayFlipping`, which represents surfaces that change the sign of one of the ray-direction components tangential to the surface at the intersection point [20]; `RayRotating`, which represents surfaces that rotate the light-ray direction by an arbitrary, but fixed, angle around the local surface normal [5]; `Refractive`, which represents standard Snell's-law refraction at the interface between optical media with different refractive indices; and `RefractiveComplex`, which represents a combination of Snell's-law refraction and rotation around the local surface normal, which can be described formally as refraction at the interface between optical media with different *complex* refractive indices [4].

Appendix E.3. Adding an existing surface-property class to interactive TIM

Surface properties are being selected in a panel described by the `SurfacePropertyPanel` class (Fig. E.15). This panel consists of a combo box that allows selection of the class of surface property, and space for editing any parameters specific to the selected surface-property class.

The following modifications of the `SurfacePropertyPanel` class add an existing surface-property class to TIM's interactive incarnation; we discuss these using the example of the surface property describing a coloured surface (`SurfaceColour` in `optics.raytrace.surfaces`).

1. Define a string constant, with a suitable name, that describes the surface property. After completion of the steps below, the contents of this string will come up as an option in the combo box for selecting a surface-property class. In our example, add the lines

```
private static final String
    SURFACE_PROPERTY_DESCRIPTION_COLOURED =
        "Coloured";
```

2. If the surface-property class requires additional parameters, define a variable that can hold the panel for editing these parameters. In our example, this panel allows separate editing of the red, green and blue (RGB) components of the surface colour, and it has a label

which helpfully points out what it is the user is editing ("Colour"). All of this can be achieved with the `LabelledDoubleColourPanel` class, using the following lines of code:

```
private LabelledDoubleColourPanel
    colourPanel;
```

In the constructor of the `SurfacePropertyPanel` class, create an instance of the panel for the additional parameters and initialise it with a default value:

```
colourPanel =
    new LabelledDoubleColourPanel("Colour");
colourPanel.setDoubleColour(
    DoubleColour.WHITE
);
```

3. The `setSurfaceProperty` method contains a chain of if statements that distinguishes between different surface-property classes of the variable `surfaceProperty`. Into this chain, link a case for the surface-property class to be added. In the code block for this if statement, add statements that make the surface-property class that currently selected in the surface-property-class combo box; set the panel for editing the surface-property-class-specific parameters to reflect the properties of the variable `surfaceProperty`; and ensure that the panel for editing the class-specific parameters is shown. The following block of code does this for our example:

```
if(surfaceProperty instanceof SurfaceColour)
{
    surfacePropertyComboBox.
        setSurfacePropertyString(
            SURFACE_PROPERTY_DESCRIPTION_COLOURED
        );
    colourPanel.setDoubleColour(
        (
            (SurfaceColour)surfaceProperty
        ).getDiffuseColour()
    );
    setOptionalParameterPanelComponent(
        colourPanel
    );
}
```

4. There is a similar chain of if statements in the `getSurfaceProperty` method. Into that chain, add a case that returns an instance of the new surface-property class with the parameters from the class-specific-parameters panel. In our case, the following lines do this:

```
if(surfacePropertyString.equals(
    SURFACE_PROPERTY_DESCRIPTION_COLOURED
))
{
    // return a shiny version of the colour
    return new SurfaceColour(
        colourPanel.getDoubleColour(),
        // specular component; white = shiny
    );
}
```

```

        DoubleColour.WHITE
    );
}

```

5. The remaining changes are to the internal class `SurfacePropertyComboBox`, which describes the combo box for selecting a surface-property class. First, the code for adding the string describing the new surface-property class to the various options available for selection by the combo box needs to be added to the constructor. This is done by adding the string constant describing the surface property defined above to the array of strings called `surfacePropertyStrings`. For our example,

```

surfacePropertyStrings.add(
    SURFACE_PROPERTY_DESCRIPTION_COLOURED
);

```

6. Finally, in the `actionPerformed` method, a case needs to be added to the chain of `if` statements which displays the class-specific-parameters panel in case the user selects the new surface-property class in the surface-property-class-selection combo box. In our example, the following code is suitable:

```

if(surfacePropertyString.equals(
    SURFACE_PROPERTY_DESCRIPTION_COLOURED)
)
{
    setOptionalParameterPanelComponent(
        colourPanel
    );
}

```

For a few classes of surface property it is necessary to edit more parameters than fit into the space reserved for the surface-property panel. The way this has been achieved in TIM is by making the surface-property-class-specific panel consist of a button which, when clicked, initiates editing of the class-specific parameters. The details of how this has been implemented can be seen by studying how the classes describing tiled and two-sided surfaces, `SurfaceTiling` and `TwoSidedSurface`, respectively, have been incorporated into the `SurfacePropertyPanel` class.

References

- [1] Hamilton, A. C. and Courtial, J., *New J. Phys.* **11** (2009) 013042.
- [2] Courtial, J., Hamilton, A. C., Šarbort, M., and Tyc, T., Natural and unnatural refraction, in preparation, 2011.
- [3] Lambert, D. et al., User guide to TIM, a ray-tracing program for forbidden optics, in preparation, 2011.
- [4] Sundar, B., Hamilton, A. C., and Courtial, J., *Opt. Lett.* **34** (2009) 374.
- [5] Hamilton, A. C., Sundar, B., Nelson, J., and Courtial, J., *J. Opt. A: Pure Appl. Opt.* **11** (2009) 085705.
- [6] Constable, G., Hamilton, A. C., and Courtial, J., Complex representation of light-ray-direction changes, in preparation, 2011.
- [7] Bryngdahl, O., *J. Opt. Soc. Am.* **64** (1974) 1092.
- [8] Berkhout, G. C. G., Lavery, M. P. J., Courtial, J., Beijersbergen, M. W., and Padgett, M. J., *Phys. Rev. Lett.* **105** (2010) 153601.
- [9] Landau, L. D. and Lifschitz, E. M., *Klassische Feldtheorie*, chapter VII, Akademie Verlag, Berlin, 1992.
- [10] Palmer, S. E., *Vision Science — photons to pheomenology*, chapter 5.3, MIT Press, 1999.
- [11] Wikipedia, Anaglyph image, http://en.wikipedia.org/wiki/Anaglyph_image.
- [12] Tyler, C. W. and Clarke, M. B., The autostereogram, in *Stereoscopic Displays and Applications*, volume 1258 of *SPIE proceedings series*, pages 182–196, Bellingham, Washington, 1990, SPIE - The International Society for Optical Engineering.
- [13] Wikipedia, Java (programming language), [http://en.wikipedia.org/wiki/Java_\(programming_language\)](http://en.wikipedia.org/wiki/Java_(programming_language)).
- [14] Wikipedia, Object-oriented programming, http://en.wikipedia.org/wiki/Object-oriented_programming.
- [15] Wikipedia, Lambertian reflectance, http://en.wikipedia.org/wiki/Lambertian_reflectance, 2010.
- [16] Wikipedia, Phong shading, http://en.wikipedia.org/wiki/Phong_shading.
- [17] Wikipedia, Ray tracing, [http://en.wikipedia.org/wiki/Ray_tracing_\(graphics\)](http://en.wikipedia.org/wiki/Ray_tracing_(graphics)).
- [18] Nelson, J., Courtial, J., and Whyte, G., *Am. J. Phys.* **76** (2008) 991.
- [19] Hicklin, J. et al., JAMA: A Java Matrix Package, <http://math.nist.gov/javanumerics/jama/>.
- [20] Hamilton, A. C. and Courtial, J., *J. Opt. A: Pure Appl. Opt.* **10** (2008) 125302.
- [21] Eclipse, <http://www.eclipse.org/>.